arXiv:1409.8327v1 [cs.SY] 29 Sep 2014

Bayesian and regularization approaches to multivariable linear system identification: the role of rank penalties

G. Prando, A. Chiuso, G. Pillonetto
Dept. of Information Engineering, University of Padova, {prandogi,chiuso,giapi}@dei.unipd.it

***Abstract*— Recent developments in linear system identification have proposed the use of non-parameteric methods, relying on regularization strategies, to handle the so-called bias/variance trade-off. This paper introduces an impulse response estimator which relies on an ℓ_2-type regularization including a rank-penalty derived using the log-det heuristic as a smooth approximation to the rank function. This allows to account for different properties of the estimated impulse response (e.g. smoothness and stability) while also penalizing high-complexity models. This also allows to account and enforce coupling between different input-output channels in MIMO systems. According to the Bayesian paradigm, the parameters defining the relative weight of the two regularization terms as well as the structure of the rank penalty are estimated optimizing the marginal likelihood. Once these hyperameters have been estimated, the impulse response estimate is available in closed form. Experiments show that the proposed method is superior to the estimator relying on the "classic" ℓ_2-regularization alone as well as those based in atomic and nuclear norm.**

I. Introduction

Linear system identification has been developed, by and large, following the so-called "parametric approach" [1], [2]. Candidate models within a prespecified model class (e.g. ARMAX/Box-Jenkins/State-space etc.) are parametrized using a finite dimensional parameter (say $\theta \in \Theta$). The most "adequate" parameter vector can be selected my minimizing a suitable cost function, most often the average squared prediction error resulting from the model (PEM). This approach heavily depends on the chosen model class and, in particular, on its complexity which will be called the *system order* hereafter.

An alternative approach which has been put forward in the recent years [3], [4], [5], [6], [7], [8], [9] is based on the theory of regularization and Bayesian statistics; factoring out minor differences, the regularization and Bayesian views just provide two alternative languages to describe the same approach. The basic idea is to choose a large enough (in principle infinite dimensional) model class so as to describe "*any*" possible system and then introduce a penalty term (regularization view) or equivalently a prior probability (Bayesian view) which is in charge of controlling the system complexity, thus facing the so called *bias-variance* dilemma.

In this paper we shall discuss, admittedly with a biased perspective due to our recent work, some Bayesian/regularization approaches which include explicit penalties for system complexity derived from the Hankel matrix built with the impulse response coefficients.

This work has been partially supported by the FIRB project "Learning meets time" (RBFR12M3AC) funded by MIUR and by the European Community's Seventh Framework Programme [FP7/2007-2013] under agreement n. 257462 HYCON2 Network of excellence.

The structure of the paper is as follows: Section II states the problem and Section III describes the classic parametric approach. In Section IV we lie the basics of regularization which are then further developed in V. Algorithmic details are provided in Section VI while numerical results are provided in Section VII.

II. Problem Formulation

We consider the following linear, causal and time-invariant (LTI) Output-Error (OE) system:

$$y(t) = G(z)u(t) + e(t) \tag{1}$$

where $y(t) = [y_1(t), .., y_p(t)]^\top \in \mathbb{R}^p$ is the p-dimensional output signal, $u(t) = [u_1(t), .., u_m(t)]^\top \in \mathbb{R}^m$ is the m-dimensional input signal, $e(t)$ is the innovation process and $G(z) := \mathcal{Z}[g](z)$ is the system transfer function. For simplicity, we will assume the presence of a delay in $G(z)$, i.e. $G(\infty) = 0$. In addition, we assume $e(t) \sim \mathcal{N}(0_p, \Sigma)$, $\Sigma = diag(\sigma_1, ..., \sigma_p)$, with $\mathcal{N}(\mu, \sigma)$ being the Gaussian distribution with mean μ and variance σ; 0_p denotes the zero-vector of size p.

The objective is to estimate the impulse response coefficients $\{g(k)\}_{k\in\mathbb{Z}^+}$ from a finite set of input-output data $\{y(t), u(t)\}_{t\in[1,N]}$.

III. Classical Methods

Classical parametric approaches assume that $G(z)$ belongs to a certain parametric class and can be completely described by a parameter vector $\theta \in \mathbb{R}^n$, with n denoting the model complexity; in this case we can adopt the notation $G_\theta(z)$. According to the prediction error methods (PEM), θ is estimated by minimizing the following function:

$$\widehat{\theta} = \arg\min_{\theta\in\mathbb{R}^n} J_P(\theta), \;\; J_P(\theta) = \sum_{t=1}^{N} \|y(t) - \hat{y}_\theta(t|t-1)\|^2 \tag{2}$$

where $\hat{y}_\theta(t|t-1) = G_\theta(z)u(t)$ denotes the one-step ahead predictor. An implicit assumption in the described procedure is that the number of available data N is much larger than the model complexity n: in particular, many interesting properties have been derived for N tending to infinity. However, PEM approaches are affected by some significant drawbacks. First, the optimization problem in (2) becomes non-convex when certain model classes are chosen, giving rise to local minima issues. Second, the selection of the model complexity n is a non-trivial step: for this purpose, many well-known

tools exist, such as cross-validation or information criteria (AIC/FPE,BIC/MDL,etc.) [1], [2] but all of them require the estimation of many models with different complexities, thus significantly increasing the computational burden. Last, and most importantly, information criteria are derived from asymptotic arguments, whose validity is of course limited when dealing with a finite number N of data. At last, the statistical properties of the obtained estimators $\widehat{\theta}$ are difficult to study [10] and experimental evidence [4], [5], [6] also shows unreliable results.

IV. Regularization approach

Recent developments in system identification have adopted regularization techniques, partly imported from the machine learning and signal processing communities, in order to overcome the described issues. In particular, regularization allows to jointly perform estimation and model selection, by moving from families of rigid, finite parametric model classes to flexible, infinite dimensional models directly described by the impulse response $\{g_\theta(k)\}_{k\in\mathbb{Z}^+}$. In this case, the estimator $\widehat{\theta}$ is obtained as the solution of the following optimization problem

$$\widehat{\theta} = \arg\min_{\theta\in\mathbb{R}^n} J_P(\theta) + J_R(\theta) \tag{3}$$

where $J_P(\theta)$ is the output data fit defined in (2), while $J_R(\theta)$ is a regularization term which penalizes certain parameters vectors θ which describe "unlikely" systems. Among the different forms of regularization $J_R(\theta)$ which have been proposed in the literature, two main classes can be identified: regularization for smoothness (aka Tikhonov regularization or Ridge regression) and regularization for selection. The first one gives rise to ℓ_2-norm penalties, while the second one arises from convex relaxations of the ℓ_0 quasi-norm (such as ℓ_1 norm or its variations like the nuclear norm) or other non-convex sparsity inducing penalties. Inspired by the approach in [11], we propose an estimator which combines these two main types of regularization by exploiting both an ℓ_2 norm penalty and a rank-penalty on the Hankel matrix of the estimated model. They are described in detail in Section IV-A and IV-B.

To simplify the derivation we consider a truncated impulse response $\{g_\theta(k)\}_{k\in\mathbb{Z}^+}$ of length T, $G_\theta(z) = \sum_{k=1}^{T} g_\theta(k)z^{-k}$, where T can always be taken large enough to catch the system dynamics. Here, $\theta \in \mathbb{R}^{Tmp}$ is the parameter vector containing all the impulse response coefficients $\{g_\theta(k)\}_{k\in[1,T]}$, $g_\theta(k) \in \mathbb{R}^{p\times m}$, with the ij-th element $[g_\theta(k)]_{ij}$ being the k-th impulse response coefficient from input j to output i:

$$\begin{aligned}\theta &= \left[\theta_{11}^\top\ \theta_{12}^\top\ \cdots\ \theta_{1m}^\top \mid \cdots \mid \theta_{p1}^\top\ \cdots\ \theta_{pm}^\top\right]^\top \\ \theta_{ij} &= \left[[g_\theta(1)]_{ij}\ [g_\theta(2)]_{ij}\ \cdots\ [g_\theta(T)]_{ij}\right]^\top\end{aligned} \tag{4}$$

for $i = 1, ..., p,\ j = 1, ..., m$. $^\top$ is the transpose operator.

We also introduce a vector notation by defining the vector of output observations, $Y \in \mathbb{R}^{Np}$,

$$Y = [y_1(1)\ \cdots\ y_1(N) \mid \cdots \mid y_p(1)\ \cdots\ y_p(N)]^\top \tag{5}$$

and the regressors matrix $\Phi \in \mathbb{R}^{Np\times Tmp}$, $\Phi = \text{blockdiag}(\phi, \cdots, \phi)$, with

$$\phi = \begin{bmatrix} \varphi_1(1) & \cdots & \varphi_m(1) \\ \vdots & \ddots & \vdots \\ \varphi_1(N) & \cdots & \varphi_m(N) \end{bmatrix}, \quad \varphi_i(j) = \begin{bmatrix} u_i(j-1) \\ u_i(j-2) \\ \vdots \\ u_i(j-T) \end{bmatrix}^\top \tag{6}$$

for $i = 1, ..., m,\ j = 1, ..., N$.
The cost function (2) can now be formulated as

$$J_P(\theta) = \|Y - \widehat{Y}_\theta\|_2^2 = \|Y - \Phi\theta\|_2^2 \tag{7}$$

with $\widehat{Y}_\theta = \Phi\theta$ being the vectorized one-step ahead predictor.

A. Regularization for smoothness and stability

This kind of regularization is derived in [4], [5] by assuming that $\{g_\theta(k)\}_{k\in\mathbb{Z}^+}$ is a realization of a zero-mean Gaussian process with autocovariance $cov(g_\theta(i), g_\theta(j)) = K(i, j)$. Since K represents a Mercer Kerner, it is associated to a unique Reproducing Kernel Hilbert Space (RKHS) $\mathscr{H}$, to which g_θ is assumed to belong. When a finite impulse response $\{g_\theta(k)\}_{k\in[1,T]}$ is considered, the norm $\|g_\theta\|_{\mathscr{H}}$ defined in $\mathscr{H}$ can be expressed through a quadratic form:

$$\|g_\theta\|_{\mathscr{H}}^2 = \theta^\top K^{-1}\theta \tag{8}$$

with $K \in \mathbb{R}^{Tmp\times Tmp}$, $[K]_{ij} = cov(g_\theta(i), g_\theta(j))$. Hence, the so-called ℓ_2-type regularization is obtained by setting $J_R(\theta) = \|g_\theta\|_{\mathscr{H}}^2 = \theta^\top K^{-1}\theta$. The structure of K can account for several properties, such as the fact that the impulse response of a linear system is an exponentially decaying function, or that it should be "smooth". See for instance [5], [6], [7] for several choices, such as "stable-spline", diagonal, diagonal/correlated, tuned/correlated kernels. The specific structure of K is defined through some hyperparameters α which can be estimated by cross-validation or by marginal-likelihood maximization (following the Empirical Bayes approach) [5], [12].

B. Regularization for complexity

Let us define the block Hankel matrix built with the impulse response coefficients of system (1), $H(\theta) \in \mathbb{R}^{pr\times mc}$:

$$H(\theta) = \begin{bmatrix} g_\theta(1) & g_\theta(2) & \cdots & g_\theta(c) \\ g_\theta(2) & g_\theta(3) & \cdots & g_\theta(c+1) \\ \vdots & \vdots & \ddots & \vdots \\ g_\theta(r) & g_\theta(r+1) & \cdots & g_\theta(r+c-1) \end{bmatrix} \tag{9}$$

A classical result from realization theory [13] shows that the rank of the block Hankel matrix $H(\theta)$ equals the McMillan degree (i.e. the complexity n) of the system, if r and c are large enough. In this work r and c are chosen such that $r + c - 1 = T$ and the matrix $H(\theta)$ is as close as possible to a square matrix.

Hence, from the identification point of view, the complexity of the estimated model can be controlled by introducing a penalty on the rank of $H(\theta)$, i.e. by defining $J_R(\theta) = \text{rank}(H(\theta))$. However, this choice of $J_R(\theta)$ makes the optimization problem (3) non-smooth and non-convex. To overcome this issue, in many previous works ([3], [14])

the rank penalty has been replaced by a penalty on its convex relaxation, i.e. the nuclear norm, by defining $J_R(\theta) = \|H(\theta)\|_*$. Recall that for a matrix X the nuclear norm is defined by $\|X\|_* = \mathrm{tr}(\sqrt{X^\top X})$. However, in this paper we adopt a non-convex approximation of the rank function, following what suggested in [11] and [15]. Recall that penalizing the rank of the Hankel matrix $H(\theta)$ is equivalent to favoring the sparsity of its singular values. A direct measure of sparsity in the components of a vector x is given by its ℓ_0 norm, $\|x\|_0$, which is equal to the number of non-zero components of x. Observing that

$$\sum_i \log|x_i| \equiv \lim_{p\to 0} \frac{1}{p}\sum_i (|x_i|^p - 1) \propto \|x\|_0 \tag{10}$$

we can approximate the ℓ_0 norm of x by its Gaussian entropy measure $\sum_i \log|x_i|$. Hence, in order to achieve sparsity in the singular values of $H(\theta)$, we define the penalty $J_R(\theta) = \log|H(\theta)H(\theta)^\top| = \sum_i \log\eta_i$, with $\{\eta_i\}_{i=[1,rp]}$ being the singular values of $H(\theta)H(\theta)^\top$. This type of penalty has been chosen since it can be reformulated in terms of a quadratic form in the vector θ (see Section V-A): in a Bayesian setting this fact is exploited to define a Gaussian prior for θ, as will be shown in Section V-B. As in [16] and [14], we also consider a weighted version of $H(\theta)$, i.e.

$$\widetilde{H}(\theta) = W_2^\top H(\theta) W_1^\top \tag{11}$$

with W_1 and W_2 chosen so that the singular values of $\widetilde{H}(\theta)$ are conditional canonical correlation coefficients. Refer to [14] for a complete derivation of W_1 and W_2. In particular, we adopted the second weighting scheme described in [14].

V. Regularization for both smoothness and complexity

The estimator we propose is based on the combination of the two regularization types described in Sections IV-A and IV-B. We define

$$J_R(\theta) = \lambda_1 \log|\widetilde{H}(\theta)\widetilde{H}(\theta)^\top| + \lambda_2 \theta^\top K^{-1}\theta \tag{12}$$

where λ_1 and λ_2 are non-negative scalar regularization parameters which control the relative weight of the two regularization terms. Details on how their value can be determined will be given in Section V-B.

Remark: Observe that in (12) we have considered the weighted Hankel matrix $\widetilde{H}(\theta) = W_2^\top H(\theta) W_1^\top$. The following description will only refer to this more general case, since the non-weighted case can be easily recovered by setting $W_1 = I_{cm}$ and $W_2 = I_{rp}$. Here, I_n denotes the identity matrix of size $n \times n$.

The impulse response coefficients contained in θ are then estimated by solving the following optimization problem:

$$\begin{aligned}\widehat{\theta} = \quad & \arg \min_{\theta\in\mathbb{R}^{Tmp}} (Y-\Phi\theta)^\top \left(\Sigma^{-1}\otimes I_N\right)(Y-\Phi\theta) \\ + \quad & \lambda_1 \log|\widetilde{H}(\theta)\widetilde{H}(\theta)^\top| + \lambda_2\theta^\top K^{-1}\theta \end{aligned} \tag{13}$$

where Y and Φ have been defined in (5) and (6), respectively. In (13) the available observations are also explicitly weighted by the inverse of the output noise variance, whose value has actually to be estimated, since it is not known a priori. In the simulations that follows its value has been set equal to the sample variance of the model obtained using only the ℓ_2 type regularization (described in Section IV-A).

Also, recall that the kernel K depends on some hyper-parameters α which we consider fixed in this setting; for instance, they can be determined by cross-validation or by marginal likelihood maximization, as detailed in [4],[5] and [6]. The simulations we performed here exploit the latter procedure.

We now show how the optimization problem (13) can be solved by a sort of block-coordinate descent algorithm.

A. Variational approximation to the log-det term

As a first step to formulate the block-coordinate descent algorithm, we need to determine a closed-form solution for the minimizer of the objective function in (13). In this regard, observe that the concave term $\log|\widetilde{H}(\theta)\widetilde{H}(\theta)^\top|$ can be expressed as the minimum of a set of upper-bounding lines [11]:

$$\log|\widetilde{H}(\theta)\widetilde{H}(\theta)^\top| = \min_{\Psi\succ 0} \mathrm{tr}\left[\widetilde{H}(\theta)\widetilde{H}(\theta)^\top\Psi^{-1}\right] + \log|\Psi| - rp \tag{14}$$

with $\Psi \in \mathbb{R}^{rp\times rp}$ being a positive definite matrix of so-called variational parameters. In addition, observe that the term $[\widetilde{H}(\theta)\widetilde{H}(\theta)^\top\Psi^{-1}]$ can be rewritten as a quadratic form in θ. Indeed, letting $\Psi^{-1} = Q = LL^T$, we have

$$\begin{aligned}\mathrm{tr}\left[\widetilde{H}(\theta)\widetilde{H}(\theta)^\top Q\right] &= \mathrm{tr}\left[L^\top\widetilde{H}(\theta)\widetilde{H}(\theta)^\top L\right] \qquad (15)\\ &= \|\mathrm{vec}(\widetilde{H}(\theta)^\top L)\|_2^2 \\ &= \|(L^\top W_2^\top \otimes W_1)\mathrm{vec}(H(\theta)^\top)\|_2^2 \\ &= \|(L^\top W_2^\top \otimes W_1)P\theta\|_2^2 \\ &= \theta^\top P^\top (W_2 Q W_2^\top \otimes W_1^\top W_1) P\theta \qquad (16)\end{aligned}$$

where $P \in \mathbb{R}^{rpcm\times Tmp}$ is the matrix which vectorizes $H(\theta)^\top$, i.e. $\mathrm{vec}\left(H(\theta)^\top\right) = P\theta$.
From (14) and (16) we can upper bound the cost in (13) and re-define

$$\begin{aligned}\widehat{\theta} = \arg\min_{\theta\in\mathbb{R}^{Tmp}} & \|\overline{Y} - \overline{\Phi}\theta\|_2^2 + \lambda_2\theta^\top K^{-1}\theta \\ & + \lambda_1\theta^\top P^\top (W_2 Q W_2^\top \otimes W_1^\top W_1) P\theta \end{aligned} \tag{17}$$

with $\overline{Y} = (\Sigma^{-1/2}\otimes I_N)Y$ and $\overline{\Phi} = (\Sigma^{-1/2}\otimes I_N)\Phi$.
For fixed λ_1, λ_2 and Q, the objective function in (17) is minimized in closed form by

$$\widehat{\theta} = \left[\overline{\Phi}^\top\overline{\Phi} + A(Q,\lambda_1,\lambda_2)\right]^{-1}\overline{\Phi}^\top\overline{Y} \tag{18}$$

$$A(Q,\lambda_1,\lambda_2) = \lambda_1 P^\top (W_2 Q W_2^\top \otimes W_1^\top W_1) P + \lambda_2 K^{-1}$$

Next section will introduce a Bayesian perspective which allows to treat λ_1, λ_2 and Q as hyper-parameters and to estimate them by marginal likelihood maximization.

B. Hyper-parameters estimation

From a Bayesian point of view, the minimizer in (18) can be viewed as the MAP estimate of θ once defined the data distribution and prior:

$$Y|\theta \sim \mathcal{N}\left(\Phi\theta, \Sigma \otimes I_N\right), \quad \theta \sim \mathcal{N}\left(0_{Tmp}, [A(Q,\lambda_1,\lambda_2)]^{-1}\right)$$

Observe that, exploiting the approximation to the log-det term described in (14) and to its reformulation as a quadratic form, it is possible to define a Gaussian prior for θ as in (16).

Within this Bayesian setting, the regularization coefficients λ_1, λ_2 and the matrix of variational parameters Q can be treated as hyper-parameters; thus, following the Empirical Bayes Paradigm, they can be estimated by maximizing the marginal likelihood:

$$\widehat{Q}, \widehat{\lambda_1}, \widehat{\lambda_2} = \arg\min_{0 \prec Q < 1, \lambda_1, \lambda_2 > 0} \mathcal{L}\left(Q, \lambda_1, \lambda_2\right) \quad (19)$$

$$\mathcal{L}\left(Q, \lambda_1, \lambda_2\right) = Y^\top \Lambda^{-1} Y + \log|\Lambda| \quad (20)$$

with $\Lambda = \Sigma \otimes I_N + \Phi\left[A(Q,\lambda_1,\lambda_2)\right]^{-1}\Phi^\top$.

Hence, once estimated $\widehat{Q}$, $\widehat{\lambda}_1$ and $\widehat{\lambda}_2$ through (19), their values can be plugged in into (18) to find the desired estimate of θ. Section VI will explain in detail how the estimation algorithm has been actually implemented.

VI. Algorithm implementation

As previously cited, the final estimate of θ is determined through a block-coordinate descent algorithm which alternatively optimizes $\widehat{\theta}$ using (18) (which can be done in closed form for fixed $\widehat{Q}$, $\widehat{\lambda}_1$, $\widehat{\lambda}_2$) and updates $\widehat{Q}$, $\widehat{\lambda}_1$, $\widehat{\lambda}_2$ through (19). Our algorithmic implementation exploits the following variant of (19) to optimize $\widehat{\lambda}_1$ and $\widehat{\lambda}_2$:

$$\widehat{\lambda}_1, \widehat{\lambda}_2 = \arg\min_{\lambda_1, \lambda_2 > 0} \mathcal{L}\left(\widehat{Q}, \lambda_1, \lambda_2\right) \quad (21)$$

after $\widehat{Q}$ has been fixed based on the current impulse response estimate (the details will be given in Section VI-A).

The iterations are stopped when the negative log likelihood does not decrease. Note that no guarantees of achieving a local minimum can be given.

Let $\widehat{\theta}^{(k)}$ denote the estimate of θ at the k-th iteration; the notations $\widehat{Q}^{(k)}$, $\widehat{\lambda}_1^{(k)}$, $\widehat{\lambda}_2^{(k)}$ will have analogue meaning. The algorithm can be summarized as follows:

1) Set K to be the kernel estimated using marginal likelihood optimization when no rank penalty is included, as done in [4].
2) Set $\widehat{\theta}^{(0)}$ equal to the estimate obtained using only the ℓ_2 type regularization.
3) Define $\widehat{Q}^{(0)}$ using the procedure in Section VI-A.
4) Determine $\widehat{\lambda}_1^{(0)}$ and $\widehat{\lambda}_2^{(0)}$ solving (21) with $\widehat{Q} = \widehat{Q}^{(0)}$.
5) At the $(k+1)$-th iteration
   - Compute $\widehat{\theta}^{(k+1)} = \left[\overline{\Phi}^\top \overline{\Phi} + A\left(\widehat{Q}^{(k)}, \widehat{\lambda}_1^{(k)}, \widehat{\lambda}_2^{(k)}\right)\right]^{-1} \overline{\Phi}^\top \overline{Y}$
   - Determine $\widehat{Q}^{(k+1)}$ as in Section VI-A.
   - Update $\widehat{\lambda}_1^{(k+1)}$ and $\widehat{\lambda}_2^{(k+1)}$ solving (21) with $\widehat{Q} = \widehat{Q}^{(k+1)}$.
   - Stop to iterate if $\mathcal{L}(\widehat{Q}^{(k+1)}, \widehat{\lambda}_1^{(k+1)}, \widehat{\lambda}_2^{(k+1)}) \geq \mathcal{L}(\widehat{Q}^{(k)}, \widehat{\lambda}_1^{(k)}, \widehat{\lambda}_2^{(k)})$ and choose $\widehat{\theta}^{(k)}$ as the final θ estimate.

Remark: Experimental evidence shows that, when optimizing w.r.t. λ_1 and λ_2, it is convenient to set a lower bound on the value of λ_2. As a result, for instance, stability of the estimator is preserved. The constraints on the hyperparameters can be seen as an hyper regularizer which limits the degrees of freedom due to hyperparameter estimation [17].

A. Update of the matrix Q

Let us consider equation (15) and let $\widetilde{H}(\theta) = U(\theta)S(\theta)V(\theta)^\top$ denote the singular value decomposition of $\widetilde{H}(\theta)$, with $S(\theta) = diag(s_1, ..., s_{pr})$. To simplify the notation, in the following we will omit the dependence of U, S and V on θ. We can rewrite (15) as follows:

$$\text{tr}\left[\widetilde{H}(\theta)\widetilde{H}(\theta)^\top Q\right] = \text{tr}\left[US^2U^\top Q\right] \quad (22)$$

Let $Q = U_Q S_Q V_Q^\top$ be the singular value decomposition of Q, with $S_Q = diag(s_1^Q, ..., s_{pr}^Q)$; if we set $U_Q = V_Q = U$, from (22) we have $\text{tr}[\widetilde{H}(\theta)\widetilde{H}(\theta)^\top Q] = \text{tr}[S^2 S_Q]$. Recalling that the term $\text{tr}[\widetilde{H}(\theta)\widetilde{H}(\theta)^\top Q]$ is included in the regularization function $J_R(\theta)$, from this last equation we can see that the singular values of Q act as penalties on the squares of the singular values of $\widetilde{H}(\theta)$: the larger the first ones, the smaller will be the estimated latter ones.

If the Hankel matrix $\widetilde{H}(\theta_0)$ of the true system was known, a natural choice for Q would be $U_Q = U_0$ and $s_i^Q = 1/s_{0,i}^2$, where $s_{0,i}$ denotes the i-th singular value of $\widetilde{H}(\theta_0)$. However, since at each iteration of the algorithm described in Section VI an impulse response estimate $\widehat{\theta}^{(k)}$ is available, we can exploit it to build $\widehat{Q}^{(k)}$. Namely, letting $\widetilde{H}(\widehat{\theta}^{(k)}) = U^{(k)}S^{(k)}V^{(k)^\top}$ with $S^{(k)} = diag(s_1^{(k)}, ..., s_{rp}^{(k)})$ being the singular value decomposition of $\widetilde{H}(\widehat{\theta}^{(k)})$, we set:

$$\widehat{Q}^{(k)} = U^{(k)}\ diag\left(1/(s_1^{(k)})^2, ..., 1/(s_{rp}^{(k)})^2\right) U^{(k)^\top} \quad (23)$$

By means of the simulations we performed, we observed that the singular values of $\widehat{Q}^{(k)}$ as defined in (23) may become very large, probably providing excessive prior along certain directions (the columns of $U^{(k)}$ which, we recall, is estimated and thus subject to uncertainty). Thus we found that a less committing prior, which just gives the same weight to all small singular values below a certain threshold, is to be preferred. In order to identify this threshold we recall that the matrix $\widetilde{H}(\widehat{\theta}^{(k)})$ has been formed using sample covariances (see Section IV-B). Thus we recall a result on the uniform convergence of sample covariances (see e.g. Theorem 5.3.2 in [18]), which shows that, under mild assumptions, the difference between the true and the estimated sample covariance is (uniformly in the lag) of the order of $O\left(\sqrt{\frac{\log(\log(N))}{N}}\right)$, where N is the number of available observations. Therefore, the first "noise" singular value of $\widetilde{H}(\widehat{\theta}^{(k)})\widetilde{H}(\widehat{\theta}^{(k)})^\top$ (i.e. that corresponding to the noise subspace, orthogonal to the image of $\widetilde{H}(\theta_0)\widetilde{H}(\theta_0)^\top$), is expected to be of the size $O\left(\frac{\log(\log(N))}{N}\right)$. Thus it is to be expected that singular values above that threshold are due to "signal" components while the smaller ones may be corrupted by noise. Hence, we redefine the singular values $s_i^{\widehat{Q}^{(k)}}$ of $\widehat{Q}^{(k)}$ as

$$s_i^{\widehat{Q}^{(k)}} = \begin{cases} \left(s_i^{(k)}\right)^{-2} & \text{if } s_i^{(k)} \geq \sqrt{c\frac{\log(\log(N))}{N}} \\ \nu(N) & \text{if } s_i^{(k)} < \sqrt{c\frac{\log(\log(N))}{N}} \end{cases} \quad (24)$$

where c is a constant which, in the simulation results, we have taken equal to the number of rows of the Hankel matrix. The saturation value $\nu(N)$ is defined as $\nu(N) := \frac{10N}{c\log(\log(N))}$. Thus, we replace the update in (23) by

$$\widehat{Q}^{(k)} = U^{(k)}\ diag\left(s_1^{\widehat{Q}^{(k)}}, ..., s_{rp}^{\widehat{Q}^{(k)}}\right) U^{(k)\top} \qquad (25)$$

VII. Numerical experiments

We now test and compare the proposed estimator on some Monte Carlo studies, generated under three scenarios, S1, S2 and S3, described below. In all cases the innovation process $e(t)$ is a zero-mean white noise with standard deviation chosen randomly at each run in order to guarantee that the signal to noise ratio on each output channel is a uniform random variable in the interval $[1, 4]$ for S1 and S2 and $[1, 10]$ for S3.

S1) We consider a fixed fourth order system with transfer function $G(z) = C(zI - A)^{-1}B$ where

$$A = \text{blockdiag}\left(\begin{bmatrix} .8 & .5 \\ -.5 & .8 \end{bmatrix}, \begin{bmatrix} .2 & .9 \\ -.9 & .2 \end{bmatrix}\right)$$

$$B = [1\ 0\ 2\ 0]^\top \quad C = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & .1 & 0 & .1 \\ 20 & 0 & 2.5 & 0 \end{bmatrix}$$

The input is generated, for each Monte Carlo run, as a low pass filtered white noise with normalized band $[0, \zeta]$ where ζ is a uniform random variable in the interval $[0.8, 1]$. $N_{MC1} = 200$ Monte Carlo runs are considered.

S2) For each Monte Carlo run $G(z)$ is generated randomly using the Matlab function `drmodel` with 3 outputs and 1 input while guaranteeing that all the poles of $G(z)$ are inside the disc of radius .85 of the complex plane. System orders are randomly chosen from 1 to 10. The input $u(t)$ is zero-mean unit variance white noise (similar behavior is obtained with low pass noise, not reported here for reasons of space). $N_{MC2} = 200$ Monte Carlo runs are considered.

S3) For each Monte Carlo run a SISO continuous-time system is generated using the Matlab function `rss`. System order is randomly chosen from 1 to 30. Each continuous-time system is sampled at 3 times the bandwith in order to derive the corresponding discrete-time system. The input $u(t)$ is zero-mean unit variance white noise filtered through a randomly generated second order filter. $N_{MC3} = 120$ Monte Carlo runs are considered.

We now consider the following algorithms:

1) ATOM: The estimator proposed in [8] which adopts a regularization based on the atomic norm of the transfer function to be estimated. The algorithm run on a set of $h \times k = 32 \times 29 = 928$ atoms built from the impulse responses of second order linear systems

$$G_{hk}(z) = C\frac{z}{(z - p_{hk})(z - p_{hk}^*)} \qquad p_{hk} = \rho_h e^{j\theta_k}$$

where $\rho_h \in [0.41 : 0.02 : 0.99\ 0.995\ 0.999] \in \mathbb{R}^{32}$
$\theta_k \in [(\pi/30) : (\pi/30) : (\pi - \pi/30)] \in \mathbb{R}^{29}$

and C is determined to guarantee unit norm. We rely on the `glmnet` package [19] for the numerical implementation. This algorithm is tested only on the SISO scenario S3.

2) PEM: The classical PEM approach, as implemented in the `pem.m` function of the MATLAB System Identification toolbox.
3) SS: The stable-spline estimator developed in [4] and [6], applied independently on each output channel. First order stable splines are used in scenario S1 and S3, while second order stable splines are adopted in S2. For both scenarios an "OE" model class is assumed.
4) SSNN: The estimator proposed in [14] which combines the ℓ_2-type penalty of Section IV-A (with kernel estimated by the SS algorithm in 2) with a nuclear norm penalty on the weighted Hankel matrix $\widetilde{H}(\theta)$. The weights W_1 and W_2 are computed as illustrated in [14]. The regularization parameters λ_1 and λ_2 are estimated through cross-validation on a predefined grid with 10 values for each hyperparameter; the "location" of the grid is chosen to optimize performance.
5) SSR: The estimator (13) obtained through the algorithm described in Section VI. Both the Hankel matrix $H(\theta)$ in (9) and its weighted version $\widetilde{H}(\theta)$ are considered and compared.

The complexity of the models estimated through PEM is the default one provided by MATLAB, while for the other algorithms the length T of the estimated impulse response was set to 80 for S1, 50 for S2 and 60 for S3. All the considered estimators are obtained using $N = 500$ pairs of input-output data for scenarios S1 and S2, while $N = 1000$ data are used for S3. Their performances are compared by evaluating the *Average Impulse Response Fit*, such defined:

$$\mathcal{F}(\widehat{\theta}) = \frac{1}{pm}\sum_{i,j} 100\left(1 - \frac{\|\theta_{ij}^0 - \widehat{\theta}_{ij}\|}{\|\theta_{ij}^0 - \bar{\theta}_{ij}^0\|}\right),\ \bar{\theta}_{ij}^0 = \frac{1}{T}\sum_{k=1}^{T} g_{ij}^0(k) \qquad (26)$$

where θ_{ij} has been defined in (4), while θ_{ij}^0 contains the true coefficients $\{g_{ij}^0(k)\}_{k=1,..,T}$ of the impulse response from input j to output i.

A. Results

The boxplots in Figure 1 prove the effectiveness of the proposed estimator: in the three experimental setups here considered, it outperforms the other pre-existing methods, even if SS gives comparable performances on S3. In particular, its implementation with the weighted version of the Hankel matrix, $\widetilde{H}(\theta)$, seems preferable to the non-weighted one.

The tests performed on S3 also show how the regularization based on the atomic norm may lead to unreliable results, compared to the other approaches here evaluated.

VIII. Conclusion and Future Work

We have presented a novel regularization-based approach for linear system identification. The ℓ_2 penalty combines two types of regularization, thus jointly enforcing "smoothness" through a "stable-spline" penalty as well as"low-complexity"

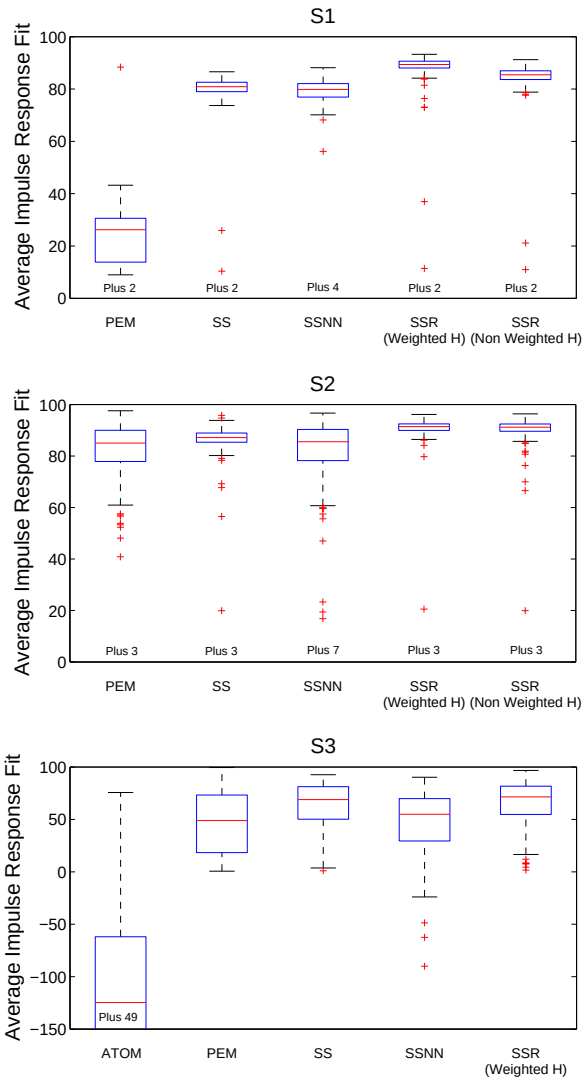

Fig. 1. Boxplot of the Average Impulse Response Fit for the three experimental scenarios S1, S2 and S3.

through a relaxation of a rank penalty on the Hankel matrix. The reformulation of the rank penalty as an ℓ_2 cost provides a Bayesian interpretation of the proposed estimator, which, in turn, can be used to perform hyperparameter selection via Marginal Likelihood maximization.

Simulation results show the performance improvements achievable through this novel method w.r.t. to other pre-existing approaches.

In our future work we plan to further investigate the properties of the kernel $A(Q, \lambda_1, \lambda_2)^{-1}$ which characterizes the new ℓ_2 penalty derived in this paper as well as to study the relation with other identification algorithms based on complexity penalties. In addition, different ways to update the matrix Q will be explored and compared to the one here proposed.

TABLE I

MEDIANS OF $\mathcal{F}(\widehat{\theta})$ IN SCENARIOS S1, S2 AND S3.

	PEM	SS	SSNN	SSR ($\widetilde{H}(\theta)$)	SSR ($H(\theta)$)
S1	24.88	80.90	79.89	89.34	85.46
S2	83.40	87.15	80.20	91.40	91.20
S3	48.94	68.93	54.88	71.53	-

REFERENCES

[1] L. Ljung, *System Identification - Theory for the User*, 2nd ed. Upper Saddle River, N.J.: Prentice-Hall, 1999.

[2] T. Söderström and P. Stoica, *System Identification*. Prentice-Hall, 1989.

[3] M. Fazel, H. Hindi, and S. P. Boyd, "A rank minimization heuristic with application to minimum order system approximation," in *In Proceedings of the 2001 American Control Conference*, 2001, pp. 4734–4739.

[4] G. Pillonetto and G. De Nicolao, "A new kernel-based approach for linear system identification," *Automatica*, vol. 46, no. 1, pp. 81–93, 2010.

[5] G. Pillonetto, A. Chiuso, and G. De Nicolao, "Prediction error identification of linear systems: a nonparametric Gaussian regression approach," *Automatica*, vol. 47, no. 2, pp. 291–305, 2011.

[6] T. Chen, H. Ohlsson, and L. Ljung, "On the estimation of transfer functions, regularizations and Gaussian processes - revisited," *Automatica*, vol. 48, no. 8, pp. 1525–1535, 2012.

[7] G. Pillonetto, F. Dinuzzo, T. Chen, G. D. Nicolao, and L. Ljung, "Kernel methods in system identification, machine learning and function estimation: a survey," *Automatica*, March, 2014.

[8] P. Shah, B. Narayan Bhaskar, G. Tang, and B. Recht, "Linear System Identification via Atomic Norm Regularization," in *Proc. of Conference on Decision and Control*, 2012, pp. 6265–6270.

[9] K. Bekiroglu, B. Yilmaz, C. Lagoa, and M. Sznaier, "Parsimonious model identification via atomic norm minimization," in *Proc. of European Control Conference*, 2014.

[10] H. Leeb and B. M. Potscher, "Model selection and inference: Facts and fiction," *Econometric Theory*, vol. 21, no. 01, pp. 21–59, 2005.

[11] D. P. Wipf, "Non-convex rank minimization via an empirical bayesian approach." N. de Freitas and K. P. Murphy, Eds. AUAI Press, 2012, pp. 914–923.

[12] A. Aravkin, J. Burke, A. Chiuso, and G. Pillonetto, "On the estimation of hyperparameters for empirical bayes estimators: Maximum marginal likelihood vs minimum mse," *Proc. of SYSID 2012*, 2012.

[13] R. Brockett, *Finite dimensional linear systems*, ser. Series in decision and control. Wiley, 1970.

[14] A. Chiuso, T. Chen, L. Ljung, and G. Pillonetto, "Regularization strategies for nonparametric system identification," in *Proc. of IEEE Conf. on Dec. and Control (CDC2013)*, 2013.

[15] K. Mohan and M. Fazel, "Iterative reweighted least-squares for matrix rank minimization," in *Proc. of Allerton Conference on Communications, Control, and Computing*, 2010.

[16] A. Hansson, Z. Liu, and L. Vandenberghe, "Subspace system identification via weighted nuclear norm optimization," *CoRR*, vol. abs/1207.0023, 2012.

[17] G. Pillonetto and A. Chiuso, "Tuning complexity in kernel-based linear system identification: The robustness of the marginal likelihood estimator," in *Proc. of European Control Conference*, 2014.

[18] E. Hannan and M. Deistler, *The Statistical Theory of Linear Systems*. Wiley, 1988.

[19] J. Qian, T. Hastie, J. Friedman, R. Tibshirani, and N. Simon, "Glmnet for matlab," $http://www.stanford.edu/\protect\kern-.1667em\relax\protect\kern-.1667em\relax\s 2013.

[20] T. Chen, M. S. Andersen, L. Ljung, A. Chiuso, and G. Pillonetto, "System identification via sparse multiple kernel-based regularization using sequential convex optimization techniques," *IEEE Transactions on Automatic Control*, 2014.

[21] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.

[22] A. Tikhonov and V. Arsenin, *Solutions of Ill-Posed Problems*. Washington, D.C.: Winston/Wiley, 1977.